

# Natural Roll Server Operation

This section covers the following topics:

- Roll Server System Requirements
  - Formatting the Roll File
  - Starting the Roll Server
  - Roll Server Messages, Condition Codes and Abend Codes
  - Return Codes and Reason Codes of the Roll Server Request
  - Operating the Roll Server
  - Roll Server Performance Tuning
  - Roll Server User Exits
  - Roll Server Restartability
- 

## Roll Server System Requirements

This section describes the roll server system requirements. The following topics are covered:

- APF Authorization
- System Linkage Index
- Virtual Storage
- CF Structure
- XCF Signalling Paths

## APF Authorization

Copy the following Roll Server load modules to an APF-authorized library:

- NATRSMAT
- NATRSSTT
- NATRSRDT
- NATRSWRT
- NATRSTRT
- NATRSDRF
- NATRSCFI

The NATRSMAT module is delivered with AC=1.

## System Linkage Index

As the Roll Server reserves one system linkage index (System LX), check whether there is a high enough value of NSYSLX in member IEASYSxx of library SYS1.PARMLIB.

If you terminate the Roll Server address space, the address space ID is no longer available. It becomes available again with the next IPL because a System LX has been used.

To avoid this, deactivate the Roll Server with the DEAL operator command and restart it afterwards.

Once a System LX has been reserved, it is reused with every restart of the Roll Server until the next IPL is done.

## Virtual Storage

ECSA	84 bytes
Private program storage	30 KB above
Fixed subpool storage (incl. ELSQA):	10 KB below, 50 KB above
LRB directory:	
100 slots per roll file	4 KB above
Every additional roll file	30 KB above
Working storage	depending on load, about 1 MB above

## CF Structure

The space required can be calculated using the following formula:

$$24 \text{ KB} + (\text{RFN} + 1) * 1 \text{ KB} + (\text{RFS} + 8) * 160 \text{ bytes}$$

where RFN denotes the number of roll files and RFS denotes the total number of roll-file slots in all roll files.

### Example:

There are five roll files with 1000 slots each.

$$24 \text{ KB} + 6\text{KB} + 5008 * 160 \text{ bytes} = 24 \text{ KB} + 6 \text{ KB} + 783 \text{ KB} = 813 \text{ KB}$$

The CF structure should thus be defined with 820 KB:

```
STRUCTURE NAME(NATROLLS) SIZE(820) PREFLIST(CF1)
```

## XCF Signalling Paths

If you use the Roll Server in a Sysplex environment, the Roll Server's instances communicate via the XCF Signalling Services. As the default XCF group name, the first eight bytes of the CF structure name are used.

If you want to specify your own XCF groupname, use the NATRSUX2 user exit. For more information on this user exit see NATRSUX2 User Exit.

## Formatting the Roll File

To format the roll file, proceed as follows:

1. Allocate it as a physical, sequential dataset with a fixed-record format.
2. Format it using module NATRSRFI.

During formatting, the roll file is converted to BDAM format; depending on the type of disk, a corresponding block size is used. To format, enter the following parameter string under the DD name RFIPARMS in module NATRSRFI:

```
function,dd-name,device-type,slot-size,number-of-blocks
```

All parameters are positional; their meaning is explained in the table below:

Parameter	Description
function	<p><b>CINIT</b> - The roll file is formatted initially.</p> <p><b>RESET</b> - All roll file slots are reset (marked as free). You can only use this parameter value if the roll file has already been formatted. Then, the only other parameter allowed is dd-name.</p>
dd-name	The name of the DD statement under which the roll file has been specified.
device-type	One of the following devices: 3380, 3390, 9345.
slot-size	<p>The size of a roll file slot in bytes. This size is rounded to the next higher multiple of the block size used.</p> <p>It is recommended to initially use a slot size equal to the size of the Natural thread. Then look at the Roll Server statistics. They also show the largest occurrence of a thread size. Use this value to reduce the slot size, if necessary.</p>
number-of-blocks	<p>The number of blocks to be allocated. If the number of blocks specified is larger than the primary space allocation, additional extents are allocated. In this way, a roll file can extend beyond several volumes. With Device Types 3380 and 3390, two blocks per track are created.</p> <p>To calculate the value of <b>number-of-blocks</b> based on the required <b>number-of-slots</b> (= maximum number of concurrent users), use the following formula:</p> <p><b>ceiling (slot-size / block-size) * number-of-slots</b></p> <p>where <b>block-size</b> depends on the device type in the following way:</p> <p>3380: 23476 bytes (half a track)  3390: 27998 bytes (half a track)  9345: 22928 bytes (FBA disk)</p>

To calculate the required disk space in cylinders (SPACE parameter of the DD statement), use the following formula:

$$\text{number-of-cylinders} = \text{ceiling} (\text{number-of-blocks} / 30)$$

This calculation bases on the fact that you have 15 tracks per cylinder, each track comprising two blocks.

**Note:** If you specify a value greater than the calculated **number-of-cylinders**, you will waste disk space.

#### Example:

```
// EXEC NATRSRFI
// RFIPARMS DD *
CINIT,ROLLF1,3390,150000,200
// ROLLF1 DD DSN = rollfilename
```

After the roll file is formatted, a WTO message is issued specifying the number of slots formatted. If an error has occurred during formatting, a WTO message is issued and formatting is terminated with Condition Code 20.

## Notes Concerning the Formatting or Resetting of Roll Files

- You can format or reset several roll files at once by specifying several parameter strings under the DD-name parameter.
- You cannot format or reset a roll file while the roll server is active.
- When you format a roll file initially, you cannot restart those Natural sessions that are still active during

formatting.

- If the Coupling Facility is used, the roll file may be formatted only if the Coupling Facility structure has been deleted using the appropriate OS/390 command.

## Starting the Roll Server

You start the Roll Server either as a batch job or as a started task by executing module NATRSMAT.

To start, enter the following parameter string under DD names ROLLF1 to ROLLF5:

```
subsystem-id,number-of-roll-files,number-of-LRB-slots,LRB-slot-size,
CF-structure-name,low-water-mark,high-water-mark,force
```

All parameters are positional and must be separated by a comma. Their meaning is explained in the table below and by the examples for starting the roll server as a batch job.

Parameter	Possible Values	Default	Comment
<i>subsystem-id</i>	4-byte non-blank string	NAT3	The specified value must match the value of the Natural profile parameter SUBSID. <b>Note:</b> With Natural under CICS, refer to the ROLLSRV parameter in the NCMDIR macro for setting the appropriate subsystem ID.
<i>number-of-roll-files</i>	1 - 5	1	
<i>number-of-LRB-slots</i>	1 - 32767	none	The number of LRB slots multiplied by the slot size must not exceed 2 GB.  The same number of LRB slots is assigned for each LRB, i.e. for each roll file used. The total number of LRB slots is calculated by the formula:  <i>number-of-roll-files * number-of-LRB-slots</i>
<i>LRB-slot-size</i>	any numeric value	roll file slot size	Value in number of bytes.
<i>CF-structure-name</i>	any valid structure name	none	If you specify less than 16 characters, blanks are added. Only specify this parameter if you use the Coupling Facility (with Parallel Sysplex).
<i>low-water-mark</i>	1 - 9	7	Specifies the low water mark in steps of ten percent of the number of LRB slots. Only specify this parameter if you do not use the CF.
<i>high-water-mark</i>	1 - 10	8	Analogous to low-water-mark parameter. Value "10" means that the staging task will never be activated. It is only recommended to specify "10" if the LRB is large enough to serve all simultaneously active Natural sessions.
<i>force</i>	F or none	none	Enables the forced restart of the Roll Server after it has been cancelled.

## Examples for Starting the Roll Server as a Batch Job

```
//EXEC NATRSMAT,PARM='N230,,1000'
```

The subsystem ID is N230, one roll file is used with 1000 slots. The slot size used is identical with the roll file's slot size. The low water mark is 70% (default), the high water mark is 80% (default).

```
//EXEC NATRSMAT,PARM=',5,1000,150000,,3,4'
```

The subsystem ID is NAT3 (default), five roll files are used with 1000 slots each. The LRB slot size is 150000 bytes. The low water mark is 30%, the high water mark is 40%.

As in the first example, but the difference is that the Coupling Facility is used and the structure NAT\_STRUCT must be defined in the CF policy.

## Roll Server Messages, Condition Codes and Abend Codes

On every important event during execution and on every error, the Roll Server issues a WTO message in the form "NATROLLS -- *text*".

The following condition codes are used:

<b>0</b>	Normal end
<b>12</b>	Wrong parameter input
<b>16</b>	Runtime error
<b>20</b>	Abend has occurred
<b>&gt;100</b>	Initialization error

## User Abend Codes

When an unexpected return code is issued by an XCF or XES Service Call, an abend with a dump is forced. Register 14 of the abend register contains the reason code. To find a description of the reason, refer to *Programming: Sysplex Services Reference* (IBM documentation). If the error was not environment-specific, send the dump to Software AG support.

The following user abend codes are used:

<b>Abend Code</b>	<b>Module</b>	<b>Cause</b>
<b>U0200</b>	NATRSMAT	IXLCONN failed
<b>U0201</b>	NATRSMAT	IXLFORCE failed
<b>U0202</b>	NATRSMAT	IXLLIST failed
<b>U0203</b>	NATRSMAT	IXLDISC failed
<b>U0204</b>	NATRSMAT	IXCLEAVE failed
<b>U0301</b>	NATRSSTT	IXLLIST failed
<b>U0302</b>	NATRSSTT	IXCMSGO failed
<b>U0401</b>	NATRSCFI	IXLLIST failed
<b>U0501</b>	NATRSDRF	IXLLIST failed

## Return Codes and Reason Codes of the Roll Server Request

These are codes that Natural may receive when the Roll Server is used during a Natural session runtime. They are output by the corresponding teleprocessing interfaces (Natural CICS or Natural IMS interface). For a list of these codes, refer to the Return Codes and Reason Codes of the Roll Server Request (in the Messages and Codes documentation).

## Operating the Roll Server

The following commands can be issued to the Roll Server via the MODIFY operator command:

Command	Description
<b>TERM</b>	Stops the Roll Server and terminates the address space. The address space ID is no longer available until the next IPL.
<b>TERM,D</b>	Analogous to TERM. Additionally, the roll file directory (in-core or in the CF) is written to the roll file. If the CF is not used, every modified LRB slot is written to the roll file.
<b>DEAL</b>	The Roll Server is stopped, but the address space is not terminated. In a de-allocated status, the Roll Server can be restarted with new parameters and the old address space ID. Since with the restart of the Roll Server, all modules including the user exits, excluding NATRSMAT, are reloaded, patches can be applied during de-allocation. Furthermore, the roll files can be reformatted in de-allocated status. If you do that, however, currently active Natural sessions are no longer restartable.
<b>DEAL,D</b>	Analogous to DEAL, but with the dump option as in TERM,D.
<b>SNAP</b>	Debugging function. The Roll Server's address space is dumped to SYSUDUMP.
<b>TRSTART</b>	Debugging function. Activates the Trace Task. If the General Trace Facility (GTF) is started and enabled for user records of Type 200, the trace record is written to the GTF.
<b>TRSTOP</b>	Deactivates the Trace Task.
<b>START,parmstring</b>	Reactivates the Roll Server with the specified parameters. You can only use this command in deactivated status.

## Roll Server Performance Tuning

As a general rule for Roll Server performance tuning, give the Roll Server a higher dispatching priority than the address spaces where Natural runs.

To find out where the weaknesses in performance are, analyze system performance with the Roll-Server Statistics function of the SYSTP utility.

When looking at Roll-Server Statistics, keep an eye especially on the following values:

- The number of direct writes.  
"Direct write" means that the Natural thread that was received was written to the roll file directly. There are two possible reasons:
  1. No LRB slot available. Increase the LRB.
  2. The compressed thread was larger than a single LRB slot. Increase the LRB slot size.
- The number of direct reads.  
"Direct read" means that the requested thread was no longer in the LRB and had to be read directly from the roll file.

If the ratio of direct reads to the total number of reads is very high in a single OS/390 system, the LRB is too small (increase it).

If the ratio of direct reads to the total number of reads is very high in a Parallel Sysplex OS/390 system, this may also mean that there are many inter-system activities, which in turn means that a Natural session changes OS/390 images quite frequently during its lifetime.

- The number of staging waits (in a single OS/390 environment).  
A "staging wait" is a situation where a write request had to wait until the Staging Task had written the LRB slot to the roll file. If the ratio of staging waits to the total number of write requests is very high, this indicates that the high and low water marks are set inappropriately or that there is a bottleneck on the roll file device/roll file channel.  
Based on experience with stress tests, the following is recommended:  
If the ratio of maximal number of active users to number of LRB slots is very small, increase the high water mark. If not, decrease the high water mark.  
The difference between high water mark and low water mark should not be larger than three (30%).  
Ideally, if the number of LRB slots is definitely larger than the maximum number of concurrent users, the high water mark should be set to 10.

## Roll Server User Exits

The roll server has two user exits.

- NATRSUX1
- NATRSUX2

A sample source module is delivered for these.

### NATRSUX1 User Exit

Specifies the roll file number to be used.

Entry calling conventions:

- Register 1 points to the parameter list as follows:  
Offset X'00': Roll Server Version ("311 ")  
Offset X'04': Number of roll files (halfword)  
Offset X'06': user ID (16 bytes)  
**Note:** You can use NATRSUX1 to simultaneously support roll server versions 2.3 and 3.1 by querying the roll server version.
- Register 13 points to a 36-fullword save area.

Return calling convention:

- Register 15 contains the number of the roll file in binary format.

#### **Note:**

If access registers are modified within this user exit, these access registers must be saved and restored on return. This user exit is called in primary addressing mode with PSW Key 8. Since it runs in cross-memory mode, no SVC except SVC 13 may be used.

### NATRSUX2 User Exit

Specifies the XCF group name to be used.

Entry calling conventions:

- Register 1 points to an 8-byte area in which the group name must be generated.
- Register 13 points to an 18-fullword save area.

As a group name default, the Roll Server will use the first 8 bytes of the CF structure name.

This user exit is called in primary mode, PSW key 8 and in task mode.

## Roll Server Restartability

When the Roll Server has been shut down and you want to make those Natural threads available that have already been received, do the following:

- If the CF is not used, stop the Roll Server with the Dump Option, which is a Roll-Server operator command. For more information on operator commands see *Operating the Roll Server*.
- If the CF is used, no further action has to be taken.